

# Designing Storage Architectures for Preservation Collections

## Designing for Large Scale

**Dave Fellingner**, Chief Scientist

[dfellinger@ddn.com](mailto:dfellinger@ddn.com)

26 September 2011

# Storage Technology

## The Good News

- Disk drives have progressed from 1TB to 4TB in 3 years.
- SATA Bus transfer rates have gone from 300MB/s to 600MB/s.
- SATA drives still operate within the 10W range.

## The Bad News

- Average building block sizes are still 12 to 24 drives in a server.
- Motherboards capable of high speed data service are in the 400W-600W range.

# Building Block Size

## If each building block is efficient will the system be efficient?

- Let's move 60 people from point A to point B using Smart Cars
- If each car gets 40M/G the resultant motorcade gets 0.67M/G
- A better choice is a bus even if it gets less than 10M/G

## Let's build a storage system with 1200 disk drives

- We can use 50 servers with 24 drives each
  - We need 30kW just for the motherboards without considering power supply efficiency
  - The drives consume 12kW so the total is 42kW

## A better choice is a bigger building block

- A dedicated dual control system running efficient state machine code can control the entire access with less than 500W added for switching
- The total is  $12\text{kW} + 1.2\text{kW} + 500\text{W} = 13.7\text{kW}$

# Architect for the Use Case

## The 50 Server System

- Data service at 100GB/s peak 50GB/s with data recovery.
- CPU to storage ratio allows for map reduce applications.
- Distributed management and external switching.

## The Dedicated Data Service Module

- Data service at 40GB/s.
- High degree of efficiency since internal data channels are less than 2W each.
- High level of consistency due to aggregation.
- Highly simplified and centralized management.
- Increased MTBF.
- 50% reduction in space requirements.
- Dynamically manageable power conservation control.
- Centralized service channels allowing on the fly data reduction.

# Summary

## **Build with “right-sized” blocks.**

- Balancing the requirements of availability, capacity, and bandwidth with the constraints of space, power, and cooling is non-trivial.
- The decision of CPU utilization per storage unit is crucial in considering a power budget.
- Select storage elements based upon usage;
  - Solid state storage for metadata and indexing
  - High speed rotating media for data mining and processing
  - Lower speed rotating media for general service

**The system size and use should dictate the scale and flexibility of the incremental module.**

# Thank You

Questions?

**DataDirect**<sup>™</sup>  
NETWORKS  
INFORMATION IN MOTION<sup>™</sup>